

“Express Mail” Mailing Label No. **EV422226228US**

**PATENT APPLICATION
ATTORNEY DOCKET NO. SUN-P9705-SPL**

5

10

**OPTICAL COMMUNICATION BETWEEN
FACE-TO-FACE SEMICONDUCTOR CHIPS**

15

Inventors: Robert J. Drost and William S. Coates

GOVERNMENT LICENSE RIGHTS

20 [0001] This invention was made with United States Government support under Contract No. NBCH020055 awarded by the Defense Advanced Research Projects Administration. The United States Government has certain rights in the invention.

25 **Related Application**

[0002] This application hereby claims priority under 35 U.S.C. 119 to U.S. Provisional Patent Application No. 60/460,104, filed on 3 April 2003, entitled, “Optical Communication for Face to Face Chips,” by inventors Robert J. Drost and William C. Coates (Attorney Docket No. SUN-P9705PSP).

30

BACKGROUND

Field of the Invention

[0003] The present invention relates to techniques for communicating between integrated circuit chips. More specifically, the present invention relates to a method and an apparatus for communicating between integrated circuit chips that are arranged face-to-face through optical signaling.

Related Art

[0004] Advances in semiconductor technology presently make it possible to integrate large-scale systems, including tens of millions of transistors, into a single semiconductor chip. Integrating such large-scale systems onto a single semiconductor chip increases the speed at which such systems can operate, because signals between system components do not have to cross chip boundaries, and are not subject to lengthy chip-to-chip propagation delays. Moreover, integrating large-scale systems onto a single semiconductor chip significantly reduces production costs, because fewer semiconductor chips are required to perform a given computational task.

[0005] Unfortunately, these advances in semiconductor technology have not been matched by corresponding advances in inter-chip communication technology. Semiconductor chips are typically integrated onto a printed circuit board that contains multiple layers of signal lines for inter-chip communication. However, signal lines on a semiconductor chip are about 100 times more densely packed than signal lines on a printed circuit board. Consequently, only a tiny fraction of the signal lines on a semiconductor chip can be routed across the printed circuit board to other chips. This problem creates a bottleneck that continues to grow as semiconductor integration densities continue to increase.

[0006] Researchers have begun to investigate alternative techniques for communicating between semiconductor chips. One promising technique involves integrating arrays of capacitive transmitter plates and receiver plates onto semiconductor chips to facilitate inter-chip communication. If a first chip is 5 situated face-to-face with a second chip so that transmitter plates on the first chip are capacitively coupled with receiver plates on the second chip, it becomes possible to transmit signals directly from the first chip to the second chip without having to route the signal through intervening signal lines within a printed circuit board. It is also possible to communicate in a similar manner through inductive 10 coupling by using wire loops to couple magnetic fields between chips.

[0007] Unfortunately, both capacitive and inductive coupling mechanisms decrease in strength by approximately the inverse of the distance between the chips. The decreased strength of the received signal reduces robustness of the communication mechanism and increases the complexity, power, and area of the 15 receiver and transmitter circuits. Additionally, capacitive and inductive coupling mechanisms suffer from cross-talk, or destructive coupling, from adjacent channels that increases with the distance between chips. It may consequently be difficult to assemble chips into a system with sufficient precision to allow capacitive and inductive coupling mechanisms to operate effectively.

20 [0008] Hence, what is needed is a method and an apparatus for communicating between semiconductor chips without the above-described problems.

SUMMARY

25 [0009] One embodiment of the present invention provides a system that communicates between a first semiconductor die and a second semiconductor die through optical signaling. During operation, the system converts an electrical

signal into an optical signal using an electrical-to-optical transducer located on a face of the first semiconductor die, wherein the first semiconductor die and the second semiconductor die are oriented face-to-face so that the optical signal generated on the first semiconductor die shines on the second semiconductor die.

5 Upon receiving the optical signal on a face of the second semiconductor die, the system converts the optical signal into a corresponding electrical signal using an optical-to-electrical transducer located on the face of the second semiconductor die.

[0010] In a variation on this embodiment, after generating the optical signal on the first semiconductor die, the system passes the optical signal through annuli located within metal layers on the first semiconductor die to focus the optical signal onto the second semiconductor die.

[0011] In a variation on this embodiment, after generating the optical signal on the first semiconductor die, the system uses a lens to focus the optical signal onto the second semiconductor die.

[0012] In a variation on this embodiment, after generating the optical signal on the first semiconductor die, the system uses a mirror to reflect the optical signal, so that the optical signal can shine on the second semiconductor die without the first semiconductor die having to be coplanar with the second semiconductor die.

[0013] In a variation on this embodiment, after generating the optical signal on the first semiconductor die, the system passes the optical signal through an interposer sandwiched between the first semiconductor die and the second semiconductor die. This interposer contains one or more waveguides that direct the optical signal, so that the optical signal shines on the second semiconductor die.

[0014] In a variation on this embodiment, the electrical-to-optical transducer is a member of a plurality of electrical-to-optical transducers located on the first semiconductor die, and the optical-to-electrical transducer is a member of a plurality of optical-to-electrical transducers located on the first semiconductor die. In this variation, a plurality of optical signals can be transmitted in parallel from the first semiconductor die to the second semiconductor die.

[0015] In a further variation, multiple spatially adjacent electrical-to-optical transducers in the plurality of electrical-to-optical transducers transmit the same signal, and electronic steering circuits in the first semiconductor die direct data to the multiple spatially adjacent electrical-to-optical transducers to correct mechanical misalignment in X , Y and Θ coordinates.

[0016] In a further variation, multiple spatially adjacent optical-to-electrical transducers in the plurality of optical-to-electrical transducers receive the same signal, and electronic steering circuits in the second semiconductor die direct data from the multiple spatially adjacent optical-to-electrical transducers to correct mechanical misalignment in X , Y and Θ coordinates.

[0017] In a variation on this embodiment, the electrical-to-optical transducer can be: a Zener diode, a light emitting diode (LED), a vertical cavity surface emitting laser (VCSEL), or an avalanche breakdown P-N diode.

[0018] In a variation on this embodiment, the optical-to-optical transducer can be: a P-N-diode photo-detector, or a P-I-N-diode photo-detector.

BRIEF DESCRIPTION OF THE FIGURES

[0019] FIG. 1 illustrates a checkerboard pattern for integrated circuit chips that communicate through face-to-face signaling in accordance with an embodiment of the present invention.

[0020] FIG. 2A presents a cross-sectional view of face-to-face chips in accordance with an embodiment of the present invention.

[0021] FIG. 2B illustrates a communication region for face-to-face chips in accordance with an embodiment of the present invention.

5 [0022] FIG. 3 illustrates six degrees of alignment between semiconductor chips in accordance with an embodiment of the present invention.

[0023] FIG. 4 illustrates how a communication channel can include multiple channel components in accordance with an embodiment of the present invention.

10 [0024] FIG. 5 illustrates an on-chip metal structure that directs an optical beam in accordance with an embodiment of the present invention.

[0025] FIG. 6 presents an offset-view of an on-chip metal structure that directs an optical beam in accordance with an embodiment of the present invention.

15 [0026] FIG. 7A illustrates line-of-sight optical paths in accordance with an embodiment of the present invention.

[0027] FIG. 7B illustrates focused optical paths in accordance with an embodiment of the present invention.

20 [0028] FIG. 7C illustrates reflected optical paths in accordance with an embodiment of the present invention.

[0029] FIG. 8 illustrates an interposer containing an array of wave guides in accordance with an embodiment of the present invention.

25 [0030] FIG. 9 presents a flow chart illustrating the process of communicating between chips through optical signaling in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION

[0031] The following description is presented to enable any person skilled in the art to make and use the invention, and is provided in the context of a particular application and its requirements. Various modifications to the disclosed 5 embodiments will be readily apparent to those skilled in the art, and the general principles defined herein may be applied to other embodiments and applications without departing from the spirit and scope of the present invention. Thus, the present invention is not limited to the embodiments shown, but is to be accorded the widest scope consistent with the principles and features disclosed herein.

10

Arrangement for Chip-to-Chip Communication

[0032] In FIG. 1, chips 101-113 are arranged to communicate with each other through face-to-face overlapping regions in their corners. In this arrangement, each chip can communicate with four neighboring chips. Note that 15 many other arrangements that facilitate face-to-face communication will be obvious to a practitioner with ordinary skill in the art.

[0033] In one embodiment of the present invention, the arrangement of chips illustrated in FIG. 1 comprises a computer system, where at least one chip, such as chip 104, contains one or more processors, and wherein neighboring 20 chips, 101, 102, 106 and 107, contain circuitry that communicates with the one or more processors in chip 104.

[0034] FIGs. 2A and 2B illustrate two chips 101 and 104 that communicate with each other through proximity communication. As is illustrated in FIG. 2A, chips 101 and 104 are positioned so that they overlap with the active 25 face of chip 101 facing the active face of chip 104. (Note that the term “active face” refers to the face of the integrated circuit that contains active circuitry.) The

region of overlap is referred to as the “communication region 202” through which transmitter and receiver circuits communicate using optical signals.

[0035] As shown in FIG. 2B, communication region 202 comprises a number of communication channels, including channel 204. One motivation for 5 proximity communication is to use modern fine-line chip lithography features to pack transmitter and receiver channels close to one another. Note that off-chip bonding and wiring methods such as wire-bonds, ball grid arrays, and circuit board traces have pitches on the order of 100 microns, whereas on-chip wiring can have pitches of less than micron. Hence, we can pack proximity communication 10 channels on a very tight pitch, for example on the order of 5-30 microns. Packing channels on a tight pitch enables high-bandwidth communication between chips, but also creates challenges in achieving the requisite alignment tolerances.

[0036] FIG. 3 illustrates six coordinates of alignment between the planar surfaces of chip *A* and chip *B*. Misalignments in the *X*, *Y*, and Θ coordinates 15 cause shifts and rotations between the chips surfaces.

[0037] It is possible to correct for mechanical misalignments (shifts or rotations) between the chip’s surfaces by subdividing the transmitter, receiver, or both transmitter and receiver for each channel as shown in FIG. 4 into an array of 20 channel components which we “call micro-pads.” Electronic steering circuits can then be used to direct data to and from the micro-pads to correct mechanical misalignments in *X*, *Y* and Θ coordinates.

[0038] However, this technique cannot correct for mechanical misalignment in *Z*, Ψ , and Φ coordinates. Misalignments in the *Z*, Ψ , and Φ coordinates cause gaps between the chips that reduced signal strength as the chips 25 become distant. Moreover, cross-talk significantly increases as the chip separation increases.

[0039] One embodiment of the present invention uses optical signaling techniques for proximity communication. In contrast to capacitive or inductive coupling techniques, optical signaling techniques can use focused or coherent light that falls off more slowly with Z-distance and causes less crosstalk. Hence, 5 optical signaling can alleviate the sensitivity to gaps, or Z-distance, between the chips.

[0040] As illustrated in FIG. 4, a channel 204 may include one or more channel components, such as channel component 402. More specifically, channel 204 may include one or more electrical-to-optical transducer elements for the 10 transmitter end of the channel, and may include one or more optical-to-electrical transducer elements for the receiver end of the channel. Example electrical-to-optical transducer elements include light-emitting diodes and vertical-cavity surface-emission lasers (VCSELs). Example optical-to-electrical transducer elements include P-N diodes or P-I-N diodes. Standard CMOS fabrication 15 technologies can be used to create light-emitting diodes and P-N diodes. However, VCSELs and P-I-N diodes require fabrication with 3-5 materials such as gallium-arsenide, or special fabrication steps in a CMOS process, respectively.

[0041] In order to increase the signal-to-noise ratio it is desirable to direct the light from the transmitter end to the receiver end of the circuit so each 20 channel's optical energy stays mostly within the channel. We refer to this directing process as "focusing" the light.

[0042] FIG. 5 illustrates an on-chip mechanism that focuses light. In FIG. 5, an annular ring 502 surrounds a light source 504. Optical energy is directed from light source 504 through a constrained path in the opening of 25 annular ring 502, which reduces transverse spreading of the optical energy. The annular ring structure can be repeated on multiple chip metallization layers to concentrate the beam further as illustrated in FIG. 6.

[0043] FIGs. 7A, 7B and 7C illustrate three off-chip mechanisms for directing light from the transmitter end to the receiver end of the channels. FIG. 7A illustrates line-of-sight optical paths, in which chips are simply aligned so that the transmitted beams of light fall on the receiving structures. FIG. 7B 5 illustrates focused optical paths, in which a lens structure 702 is used to focus the transmitted light. This focused optical path mechanism is more complex, but compensates for spreading of the beam of light. FIG. 7C illustrates a reflected optical path mechanism, which uses a reflector, such as mirror 704, to permit the transmitter end and receiver ends of the channel to be non-coplanar. Note that the 10 reflected path mechanism can be combined with the focused path method.

[0044] Finally, FIG. 8 illustrates another off-chip mechanism that uses an interposer 802, which contains an array of embedded waveguides to direct light from the transmitter end to the receiver end of the channel. A typical optical waveguide includes an optically transparent material with two or more indices of 15 refractivity. The channel in a waveguide has a higher index of refraction than the surrounding cladding material, so that light in the channel reflects off the walls of the waveguide and remains contained in the channel.

[0045] FIG. 9 presents a flow chart illustrating the process of communicating between chips through optical signaling in accordance with an 20 embodiment of the present invention. The process starts when an electrical signal is converted into an optical signal through an electrical-to-optical transducer located on a first semiconductor die (step 902). Next, the optical signal is directed onto a second semiconductor die using any of the techniques described above with respect to FIG. 5, FIG. 6, FIGs. 7A-7C and FIG. 8 (step 904). Finally, the optical 25 signal is converted into a corresponding electrical signal through an optical-to-electrical transducer located on the second semiconductor die (step 906).

[0046] The foregoing descriptions of embodiments of the present invention have been presented only for purposes of illustration and description. They are not intended to be exhaustive or to limit the present invention to the forms disclosed. Accordingly, many modifications and variations will be apparent to practitioners skilled in the art. Additionally, the above disclosure is not intended to limit the present invention. The scope of the present invention is defined by the appended claims.